

# Non-linear eigensolver-based alternative to traditional SCF methods

B. Gavin<sup>1</sup> and E. Polizzi<sup>1</sup>

Department of Electrical and Computer Engineering, 100 Natural Resources Road, Marcus 201,  
University of Massachusetts, Amherst, USA

(Dated: 30 January 2013)

The self-consistent procedure in electronic structure calculations is revisited using a new highly efficient and robust algorithm for solving the non-linear eigenvector problem i.e.  $H(\{\psi\})\psi = E\psi$ . This new scheme is derived from a generalization of the FEAST eigenvalue algorithm to account for the non-linearity of the Hamiltonian with the occupied eigenvectors. Using a series of numerical examples and the DFT-Kohn/Sham model, it will be shown that our approach can outperform the traditional SCF mixing-scheme techniques by providing a high converge rate, convergence to the correct solution regardless of the choice of the initial guess, and a significant reduction of the eigenvalue solve time in simulations.

Keywords: diagonalization technique, electronic structure calculations, SCF, FEAST, Pulay, DIIS

## I. INTRODUCTION

Although first-principle calculations in general, and DFT in particular, have provided a practical (i.e. numerical tractable) path for solving the electronic structure problem, they have introduced new numerical challenges on their own. Within the single electron picture, the Hamiltonian operator depends on the occupied eigenfunctions and the resulting eigenvalue problem becomes fully non-linear (i.e.  $H(\{\psi\})\psi = E\psi$ ). In practice, this *non-linear eigenvector problem* is commonly addressed using a self-consistent field method (SCF) wherein a series of linear eigenvalue problems (i.e.  $H\psi = E\psi$ ), needs to be solved iteratively until convergence<sup>1</sup>. Successfully reaching convergence by performing SCF iterations is of paramount importance to first-principle quantum chemistry and solid-state physics simulations software. A typical self-consistent iteration procedure for the discretized DFT/Kohn-Sham problem is represented in Figure 1.

Traditional SCF mixing methods employ successive approximation iterates of a fixed point mapping to generate the new input electron density at each cycle. Examples of such methods include Newton-Broyden, or other Anderson technique, and Pulay mixing techniques using direct inversion of the iterative subspace (DIIS)<sup>2-6</sup>. Using DIIS, the new input electron density at iteration  $k$ , is generated from a linear combination of a series of previous trial densities (i.e.  $\{\mathbf{n}^{(k-1)}, \mathbf{n}^{(k-2)}, \mathbf{n}^{(k-3)}, \dots\}$ ) by minimizing the successive residual errors between input and output densities.

Three main numerical difficulties arise from standard SCF procedures: (i) the linear eigenvalue problem needs to be solved repeatedly a large number of times; (ii) the robustness of the self-consistent iterations is very sensitive to the choice of the initial guess; and (iii) there is, as of yet, no robust and efficient general-purpose iterative numerical scheme for addressing the non-linear coupling with guaranteed convergence. Although much progress has been made to improve the converge rate of SCF techniques<sup>7-12</sup>, the iterations can still be found to converge very slowly or unreliably<sup>13</sup>.

In this paper, we present an efficient alternative to tra-

---

---

### Basic SCF iteration procedure

---

**Input:** number of occupied states  $M$

#### 1- Initialization

Derive an initial guess for the electron density  $\mathbf{n}$

#### 2- Hamiltonian Construction $\mathbf{H}[\mathbf{n}]$

requires also solving the Poisson equation

#### 3- Linear Eigenvalue Problem $\mathbf{H}\mathbf{x} = \mathbf{E}\mathbf{S}\mathbf{x}$

solve for the  $M$  lowest eigenpairs ( $\{E_m, \mathbf{x}_m\}$ )

#### 4- Compute electron density

$\mathbf{n} = 2 \sum_{m=1}^M |\mathbf{x}_m|^2$ ; check convergence of  $\mathbf{n}$   
continue iterations if needed

#### 5- Generate new input density for next iteration

Use mixing approach to generate new density,  
go back to step 2

**Output:** Ground state electron density  $\mathbf{n}$

---

---

FIG. 1. Basic self-consistent procedure for solving the (discretized) non-linear eigenvector problem  $\mathbf{H}[\mathbf{n}]\mathbf{x} = \mathbf{E}\mathbf{S}\mathbf{x}$  with  $\mathbf{H}$  real symmetric and  $\mathbf{S}$  symmetric positive definite ( $\mathbf{S} \neq \mathbf{I}$  using non-orthogonal basis functions), and obtaining the ground state electron density  $\mathbf{n} = 2 \sum_m |\mathbf{x}_m|^2$  (with a factor 2 for spin). For the DFT/Kohn-Sham problem, the Hamiltonian  $\mathbf{H}[\mathbf{n}] = -\Delta + \mathbf{V}_H[\mathbf{n}] + \mathbf{V}_{XC}[\mathbf{n}] + \mathbf{V}_{ext}$ , is composed of the Hartree potential  $\mathbf{V}_H$  (solution of Poisson equation), the exchange-correlation potential  $\mathbf{V}_{XC}$ , and other external potential  $\mathbf{V}_{ext}$  including the ionic potential.

ditional SCF iteration techniques that is ideally suited to address the aforementioned numerical difficulties when solving the discretized non-linear eigenvector problem  $\mathbf{H}[\mathbf{n}]\mathbf{x} = \mathbf{E}\mathbf{S}\mathbf{x}$ . The SCF problem is fully revisited using a general-purpose numerical strategy derived from a modification of the FEAST eigenvalue algorithm<sup>14</sup>. We show that the new approach, named NLFEAST (for Non-Linear FEAST), is able to outperform the traditional SCF mixing-scheme techniques. NLFEAST exhibits a

high converge rate, converges to the correct solution regardless of the choice of the initial guess, and significantly reduces the eigenvalue solve time in simulations.

The outline of this paper is as follows: In section II we briefly summarize the numerical steps and the main properties of the FEAST algorithm presented in Ref. 14 for solving the linear eigenvalue problem. In section III we present the generalization of FEAST for solving the non-linear eigenvector problem. Numerical results and capabilities of the new NLFEAST approach are then presented and discussed in Section IV.

## II. THE FEAST ALGORITHM

Within the SCF procedure in Figure 1, solving the linear and symmetric eigenvalue problem at a given iteration step becomes the most time-consuming part of the electronic structure calculations. In such computations, one can identify two main challenges: (i) discretization techniques that accommodate atomistic systems with a high level of accuracy can lead to large size system matrices  $\{\mathbf{H}, \mathbf{S}\}$ , and (ii) the number of eigenpairs needed to compute the electron density is proportional to the number of atoms in the system. In order to characterize large-scale nanostructures and complex systems of current technological interest, many thousands of eigenpairs are indeed needed. In this regard, progress in large-scale electronic structure calculations can be tied together with advances in numerical algorithms for addressing the eigenvalue problem.

The recent FEAST algorithm<sup>14</sup>, which uniquely combines accuracy, robustness, high-performance and (linear) parallel scalability, is ideally suited for addressing the electronic structure problem. FEAST is a general algorithm that can be used for solving the linear eigenvalue problem to obtain all the eigenvalues and eigenvectors within a given search interval (e.g.  $[E_{\min}, E_{\max}]$ ). FEAST's main computational tasks consist of solving a small number of independent linear systems with multiple right-hand sides along a complex contour and one reduced dense eigenvalue problem that is orders of magnitude smaller than the original one (the size of this reduced problem is of the order of the number of eigenpairs inside the search interval). A full and detailed accounting of the implementation of the FEAST algorithm is given in Ref. 14, and the basic procedure is briefly summarized in Figure 2.

Worthy of particular note are Steps 2 and 4(c) which, combined, generate a subspace  $\mathbf{Q}$  that spans the eigenvectors solutions in the user-defined interval. This is done by multiplying a given approximate eigenvector solution at iteration  $k$  by the density matrix  $\rho$  i.e.

$$\mathbf{Q}_{N \times M_0}^{k+1} = \rho \mathbf{S} \mathbf{X}_{N \times M_0}^k, \quad (1)$$

where the new subspace  $\mathbf{Q}^{k+1}$  is then obtained at the FEAST iteration  $k + 1$ . Denoting by

---

### Basic FEAST Algorithm

---

**Input:** interval  $[E_{\min}, E_{\max}]$ , and an (over)-estimation  $M_0$  of the number of eigenpairs

#### 1- Initialization

Select  $M_0 > M$  random vectors  $\mathbf{Y} \in \mathbb{R}^{N \times M_0}$

#### 2- Contour Integration

Compute  $\mathbf{Q}_{N \times M_0} = -\frac{1}{2\pi i} \int_{\mathcal{C}} dZ \mathbf{G}(Z) \mathbf{Y}_{N \times M_0}$ ,

#### 3- Rayleigh-Ritz

Form  $\mathbf{H}_{\mathbf{Q} M_0 \times M_0} = \mathbf{Q}^T \mathbf{H} \mathbf{Q}$  and  $\mathbf{S}_{\mathbf{Q} M_0 \times M_0} = \mathbf{Q}^T \mathbf{S} \mathbf{Q}$

Solve the reduced eigenvalue problem  $\mathbf{H}_{\mathbf{Q}} \Phi = \epsilon \mathbf{S}_{\mathbf{Q}} \Phi$

#### 4- Subspace Iteration

(a) Set  $E_m = \epsilon_m$  and  $\mathbf{X}_{N \times M_0} = \mathbf{Q}_{N \times M_0} \Phi_{M_0 \times M_0}$

(b) Check convergence (i.e.  $\text{Trace}\{E_m\}$ )

(c) If needed go back to step 2 with  $\mathbf{Y} = \mathbf{S} \mathbf{X}$

**Output:** All the  $M < M_0$  eigenpairs  $(\{E_m, \mathbf{x}_m\})$ .

---

FIG. 2. Basic FEAST procedure for solving the generalized eigenvalue problem  $\mathbf{H}\mathbf{x} = E\mathbf{S}\mathbf{x}$  of size  $N$  with  $\mathbf{H}$  real symmetric and  $\mathbf{S}$  symmetric positive definite (spd), and obtaining all the  $M$  eigenpairs within a given interval  $[E_{\min}, E_{\max}]$ . The density matrix appears implicitly in Step-2, using the complex contour integration of the Green's function  $\mathbf{G}(Z) = (Z\mathbf{S} - \mathbf{H})^{-1}$ . In practice, the vectors  $\mathbf{Q}$  in Step-2 can be computed using a high-order numerical integration such as Gauss-Legendre quadrature, where only a small number of linear systems,  $(Z_e \mathbf{S} - \mathbf{H})\mathbf{Q}_e = \mathbf{Y}$ , need to be solved, one for each of a number of specific Gauss nodes  $Z_e$  (associated with the weights  $\omega_e$ ) along a complex contour  $\mathcal{C}$ , i.e.  $\mathbf{Q} := \sum_{Z_e \in \mathcal{C}} \omega_e \mathbf{Q}_e$ .

$\mathbf{X}_{N \times M} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_M\}$  the  $M$  eigenvectors within the search interval, we note that the density matrix is formally given by  $\rho = \mathbf{X} \mathbf{X}^T$ . In practice, the spectral projection (1) is conveniently obtained by numerical integration of a set of solutions of independent linear systems defined along a complex contour. It can be shown that FEAST consists of a reformulation of the subspace iteration technique for spectral projectors<sup>15</sup>, which leads to an extremely robust and accurate numerical procedure. By this means the linear FEAST algorithm achieves very high convergence rate of the eigenvectors subspace, and allows the non-linear extension of FEAST (i.e. NLFEAST) to achieve the kind of performance that will be demonstrated in this paper.

The FEAST algorithm holds all the following important intrinsic properties:

- using a high-order Gauss-Legendre quadrature, 8 to 16 contour points suffice for FEAST to consistently converge in  $\sim 3$  iterations to obtain up to thousands of eigenpairs with machine accuracy;
- all multiplicities are naturally captured;

- no (explicit) orthogonalization procedure is required;
- pre-computed subspace can be reused as suitable initial guess (i.e. Step-1 in Figure 2) for solving a series of eigenvalue problems that are close one another. One can ideally take advantage of this feature for addressing, in particular: (i) Step-3 of the SCF procedure in Figure 1 along the iterations, (ii) bandstructure calculations<sup>14</sup> along the  $k$  space, or (iii) time-dependent propagation using spectral decomposition along the time-domain<sup>16</sup>;
- the inner linear system at each contour point can be solved using either direct or iterative methods;
- the algorithm can be extended efficiently for solving non-Hermitian problems (e.g. for performing complex bandstructure calculations<sup>17</sup>);
- efficient parallel implementations can be naturally addressed at three different levels: (i) many search intervals can be run independently (no overlap), (ii) each linear system can be solved independently along the complex contour (e.g. simultaneously on different compute nodes), and (iii) the linear systems can be solved in parallel (the multiple right sides can be parallelized as well). Since the search intervals can be arbitrarily narrowed within a parallel environment (i.e. increasing the ratio  $N/M_0$ ), the algorithm complexity is directly dependent on solving a single linear system of size  $N$  with  $M_0$  right-hand-sides (i.e.  $(Z_e S - H)Q_e = Y$  in Figure 2).

### III. DIRECT SOLUTION OF THE NON-LINEAR EIGENVECTOR PROBLEM

Here, we propose to further extend the capabilities of the FEAST algorithm in order to solve nonlinear eigenvector problems of the form

$$H[n]x_i = E_i S x_i, \quad n = 2 \sum_i |x_i|^2, \quad (2)$$

so that it may be used as a direct and efficient alternative to SCF methods conventionally based on a ‘‘Schrödinger-Poisson’’ iterative procedure with density mixing schemes.

Recently Yang et al.<sup>18</sup> have proposed an alternative approach to traditional SCF iterations using a direct constrained minimization (DCM) algorithm for solving the nonlinear problem (2). DCM consists of constructing and updating new search directions for the eigenvector subspace by solving a much smaller optimization problem which takes the form of a reduced nonlinear eigenvector problem. These results have motivated this current work, since a reduced eigenvalue problem can also be obtained from the FEAST algorithm, which also provides an optimal framework for performing subspace iterations. Indeed, once the search subspace  $Q$  is obtained

by the contour integration in Step-2 of Figure 2 (from a given Hamiltonian  $H[n]$ ), the resulting reduced problem in Step-3 can be expressed in a non-linear form i.e.

$$H_Q[n]\Phi = \epsilon S_Q \Phi, \quad \text{with} \quad n = 2 \sum_i |Q\Phi_i|^2, \quad (3)$$

with  $H_Q[n] = Q^T H[n] Q$ . Using a DFT/Kohn-Sham model,  $H_Q$  accounts for the projection of the non-linear Hartree and exchange-correlation terms onto the working search subspace  $Q$ . One of the primary difference between linear FEAST and the proposed non-linear FEAST algorithm, named NLFEAST, then lies in Step-3 with the formation and solution of the reduced non-linear eigenvector problem. This latter, in turn, can be solved using the traditional SCF procedure presented in Figure 1. Since the non-linearity appears at the level of the reduced system alone, more robust non-linear schemes that would have been too expensive to use on the large size original system (2), can potentially be considered for addressing the reduced system (3) (e.g. specific Newton-Raphson method<sup>19</sup>).

As in the linear case, once the solution of the reduced system is obtained, the subspace is updated and the FEAST algorithm can come back on track and perform subspace iterations up until convergence. As it will be demonstrated in our experimental results in Section IV, however, the algorithm essentially converges faster when the size of the subspace is increased with each FEAST subspace iteration. These results suggest that the size of the search subspace should not be made static for the non-linear problem. This then constitutes a major and innovative difference with the original FEAST algorithm that aims at improving the robustness of the subspace iteration technique. The Raleigh-Ritz procedure (i.e. Step-3) then uses a new search subspace  $\hat{Q}$  which consists not only of the subspace most recently produced by solution of the contour integration at step  $k$ , i.e.  $Q^{(k)}$ , but also of all (or some) of the subspaces generated in previous iterations such that  $\hat{Q} = \{Q^{(k)}, Q^{(k-1)}, Q^{(k-2)}, \dots\}$ . The subspace provided by the contour integration can be appended to  $\hat{Q}$  indefinitely until convergence, or this can be done such that only a finite number (typically 3 or 4) of the most recently generated subspaces are retained. Interestingly, this procedure is reminiscent of the one used in DIIS where new electron densities are constructed from a subspace consisting of previous generated electron densities and density residuals. However, traditional mixing schemes act on the electron density alone, and it is not always guaranteed that the density can be expressed as a linear combination of previously generated densities. In contrast, the proposed approach is expected to be much more robust by essence since a larger and refined search subspace containing all the eigenvectors of interests, is very likely to span the solution of the non-linear problem.

The resulting NLFEAST algorithm is given in Figure 3. We note that a singular value decomposition (SVD) also needs to be performed on the search subspace  $\hat{Q}$

---

**Non-linear eigenvector FEAST algorithm**


---

**Input:** interval  $[E_{\min}, E_{\max}]$  including  $M$  lowest occupied states and an initial search subspace of size  $M_0 > M$

**1- Initialization**

- a-Define an initial guess for the electron density  $\mathbf{n}$
- b-Construct Hamiltonian  $\mathbf{H}[\mathbf{n}]$

c-Select  $M_0 > M$  random vectors  $\mathbf{Y} \in \mathbb{R}^{N \times M_0}$

**2- Contour Integration at FEAST iteration  $k$** 

Compute  $\mathbf{Q}_{N \times M_0}^{(k)} = -\frac{1}{2\pi i} \int_{\mathcal{C}} dZ (Z\mathbf{S} - \mathbf{H}[\mathbf{n}])^{-1} \mathbf{Y}_{N \times M_0}$ ,

**3- Non-linear Rayleigh-Ritz**

a-Construct the subspace  $\hat{\mathbf{Q}} = \{\mathbf{Q}^{(k)}, \mathbf{Q}^{(k-1)}, \mathbf{Q}^{(k-2)}, \dots\}$

b-Compute  $\mathbf{Q} := \mathbf{U}$  using SVD i.e.  $\hat{\mathbf{Q}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$

c- SCF Procedure

- construct Hamiltonian  $\mathbf{H}[\mathbf{n}]$   
(requires also solving the Poisson equation)
- form  $\mathbf{H}_{\mathbf{Q}_{M_0 \times M_0}} = \mathbf{Q}^T \mathbf{H}[\mathbf{n}] \mathbf{Q}$  and  $\mathbf{S}_{\mathbf{Q}_{M_0 \times M_0}} = \mathbf{Q}^T \mathbf{S} \mathbf{Q}$
- solve reduced eigenvalue problem  $\mathbf{H}_{\mathbf{Q}} \Phi = \epsilon \mathbf{S}_{\mathbf{Q}} \Phi$
- compute electron density  $\mathbf{n} = 2 \sum_{m=1}^M |\mathbf{Q} \Phi_m|^2$
- Check convergence of  $\mathbf{n}$ , go back to step-c if needed with a new generated input density

**4- Subspace Iteration**

Set  $E_m = \epsilon_m$  and compute  $\mathbf{X}_{N \times M_0} = \mathbf{Q}_{N \times M_0} \Phi_{M_0 \times M_0}$

Check convergence (i.e. Trace  $\sum_m E_m \in [E_{\min}, E_{\max}]$ ).

If needed go back to step 2 with  $\mathbf{Y} = \mathbf{S} \mathbf{X}$

**Output:** All the  $M < M_0$  eigenpairs  $(\{E_m, \mathbf{x}_m\})$ .

---

FIG. 3. FEAST procedure for solving the non-linear eigenvector problem  $\mathbf{H}[\mathbf{n}]\mathbf{x} = E\mathbf{S}\mathbf{x}$  of size  $N$  with  $\mathbf{H}$  real symmetric and  $\mathbf{S}$  symmetric positive definite (spd), and obtaining all the  $M$  lowest occupied states within a given interval  $[E_{\min}, E_{\max}]$  and hence the ground state electron density. Here, Step 3-c of the algorithm is making use of a traditional SCF procedure for solving the non-linear reduced system. We note that in Step-3a, the column vectors of the newly generated subspace are appended to the matrix of the old subspace, increasing the total subspace size by  $M_0$ . This can be done such that only a finite number of the most recently generated subspaces are retained, or the subspace size can be increased until convergence. Additionally, a singular value decomposition (SVD) is performed on  $\hat{\mathbf{Q}}$  (i.e. Step 3-b), such that only the left singular vectors  $\mathbf{U}$  are used in the Rayleigh-Ritz procedure.

(i.e. Step 3-b), and only the left-singular vectors are retained. In contrast to the linear FEAST algorithm, one does not want to truncate the size of the search subspace for constructing a positive definite reduced matrix  $\mathbf{S}_{\mathbf{Q}}$ . This latter is by definition rank deficient and SVD cannot then be avoided. In our experiments thus far this step has proven to be computationally inexpensive rela-

tive to other parts of the algorithm, and even so, it can be shown that the SVD of  $\hat{\mathbf{Q}}$  can be updated at each successive FEAST iteration with complexity similar to solving a single reduced eigenvalue problem in Step-3c<sup>20</sup>. Finally, like other schemes for solving nonlinear eigenvector equations, NLFEAST requires an initial guess for the density  $\mathbf{n}$  (Step-1a of Figure 3). However, the quality of the initial guess is not important as far as achieving convergence is concerned, as it will be shown in Section IV.

In summary, NLFEAST can be seen as an inversion of the usual process of solving the nonlinear eigenvector problem (2). Using SCF-DIIS, one iteration corresponds to a traditional SCF iteration as presented in Figure 1 which requires, in particular, solving a large linear eigenvalue problem as well as the calculation of the Hartree potential by solving Poisson's equation. This scheme presents then two iterative procedures: one outer associated with the SCF, and one inner due to the eigenvalue solver (using a "black-box" eigenvalue solver, those inner iterations -intrinsic to eigenvalue algorithms- are most often hidden to the user). As discussed earlier, this linear eigenvalue problem can ideally be solved using the linear FEAST algorithm (in Figure 2) which converges in very few iterations of contour integration (typically  $\sim 3$ ). Using NLFEAST, however, those two iterative procedures happen now in reverse order. The outer iterations are intrinsic to the FEAST algorithm (i.e. subspace iterations), while the inner iterations are used for solving the non-linear reduced system (3). In addition to offering a more robust mathematical approach, NLFEAST is expected to considerably reduce the eigenvalue solve time since: (i) the reduced non-linear problem (3) is orders of magnitude smaller in size as compared to the original one (2), and (ii) the reduced problem does not need to be solved very accurately to guarantee the convergence of the outer-iterations.

**IV. NUMERICAL EXPERIMENTS AND CAPABILITIES**

In this section, we propose to demonstrate the numerical efficiency of the proposed NLFEAST algorithm for the non-linear eigenvector problem (2). DFT/Kohn-Sham/LDA calculations have been performed on various molecules using our in-house all-electron simulation framework that uses a real-space cubic finite element discretization. Some details of our modeling and real-space discretization setup have been provided in Ref. 21. The Pulay-DIIS technique has been used for performing the traditional SCF approach defined in Figure 1, as well as for solving the FEAST non-linear reduced system in Step-3c of Figure 3. Both DIIS procedures make use of a subspace composed of five successive generations of the electron densities, but the maximum number of DIIS iterations for solving the reduced system in NLFEAST has been fixed to three (i.e. the non-linear reduced system is then solved only approximately). For most of the exam-

ples, indeed, increasing the number of the inner iterations further has had no effect on the overall NLFEAST convergence rate.

### A. Performance comparisons

In all our experiments conducted thus far, including various molecules from  $H_2$  to  $C_{60}$ , the NLFEAST algorithm has outperformed SCF-DIIS both in terms of convergence rate and execution time. Three representative examples are shown in Figure 4 for the Silane, Benzene and Caffeine molecules. The relative error on




FIG. 4. Results of numerical experiments comparing the performance of our algorithm to that of DIIS for the Silane  $SiH_4$ , Benzene  $C_6H_6$  and caffeine  $C_8H_{10}N_4O_2$  molecules. The relative error on the total energy is used here as the measure of convergence. The convergence criteria for NLFEAST has been set to  $10^{-10}$  while DIIS-SCF was stopped after a maximum of 10 iterations for Silane and 15 iterations for both Benzene and Caffeine. The meaning of outer-iterations is different for both algorithms.

the total energy at each iteration is used as the measure of convergence for both approaches, as this is the most directly comparable measure between DIIS-SCF and NLFEAST. As discussed in Section III, the meaning of the outer-iterations is different for both algorithm since, for NLFEAST, it directly represents the number of contour integrations and non-linear reduced systems.

Since DIIS-SCF also takes advantage of the linear FEAST algorithm, it is possible to have a direct comparison between the number of operations counts of DIIS-SCF and NLFEAST. In both cases the contour integration is the most time consuming single numerical operation. In the three examples provided in Figure 4, for example, SCF-DIIS requires  $5\times$  to  $7\times$  more contour integrations than NLFEAST for obtaining the same level of accuracy (i.e.  $\sim 10^{-7}$ ). On the other hand, NLFEAST requires the solution of  $\sim 3\times$  as many Poisson equations as does SCF-DIIS. However, since the Poisson system

matrix only needs to be factorized once at the beginning of the calculation, the remaining Poisson solve operations become negligible in comparison to both factorizing and solving the complex linear systems that arise from the contour integration. The rest of the numerical operations involved in NLFEAST, such as solving the reduced eigenvalue problem or performing the SVD in Step-3c of Figure 3, bring no significant computational overhead provided the ratio  $N/M_0$  stays relatively large for a given search interval. In practice, it is also reasonable to use an original subspace size  $M_0$  that is typically no larger than  $\sim 500$ , if we assume that the  $\hat{Q}$  subspace keeps increasing indefinitely and that convergence is reached in less than 10 iterations. Since it is recommended for FEAST to use  $M_0 \sim 1.5\times$  the number of eigenstates<sup>14</sup>, the basic NLFEAST scheme is expected to provide good performances up to  $\sim 350$  states (i.e. 700 electrons with a factor 2 for spin). It should be noted that the case of a smaller  $N/M_0$  ratio with a relatively large  $M_0$ , can lead to degradation of performances and memory limitations only if the algorithm runs sequentially. The FEAST algorithm can indeed be readily parallelized and the applicability of NLFEAST scheme for addressing large-scale molecular systems will be discussed in Section V.

Table I summarizes the convergence results we have obtained using NLFEAST for a sample of molecules ranging in size from  $H_2$  to  $C_{60}$ , with results on the total energy which are compared with the NWChem software<sup>22</sup>. Although we did no attempt to optimize or refine our cubic finite element discretization, the total energy results are found to be in good agreement with NWChem which is making use of completely different numerical approaches and convergence criteria (for these reasons also, a direct and useful comparison of the convergence rate cannot be achieved here). It is also important to note that in the case of the non-linear eigenvector problem (2), the relative error on the total energy is not a natural criteria for measuring the convergence. For NLFEAST, the criteria is instead defined in terms of the error on the trace of the eigenvalues. From our numerical experiments, once this convergence criteria is satisfied, it also provides very low non-linear residual. For the simulations reported in Table I, we have also used: (i) a  $\hat{Q}$  search subspace that keeps increasing indefinitely until convergence, (ii) 16 nodes for the Gauss quadrature along the complex contour, (iii) a conventional initial guess as the starting point for the electron density (e.g. the all-electron result for the single atoms). In the following subsections, the robustness of NLFEAST will be discussed and demonstrated further by separately addressing each one of those points.

### B. Convergence Rate and $\hat{Q}$ subspace size

The size of the search subspace  $\hat{Q}$  (in step-3a of Figure 3) may be limited by only retaining a finite number of the most recent subspaces, or it may be extended indefinitely until convergence. For the numerical exper-

	#electrons	#iterations	$E_{\text{tot}}(\text{eV})$	NWChem
H <sub>2</sub>	2	5	-30.962	-30.959
CH <sub>4</sub>	10	5	-1091.45	-1091.69
H <sub>2</sub> O	10	7	-2065.24	-2065.48
CO	14	6	-3060.12	-3060.45
SiH <sub>4</sub>	18	9	-7907.42	-7909.78
Na <sub>2</sub>	22	8	-8785.57	-8786.61
C <sub>6</sub> H <sub>6</sub>	42	6	-6262.31	-6263.65
C <sub>8</sub> H <sub>10</sub> N <sub>4</sub> O <sub>2</sub>	102	9	-18364.5	-18360.3
C <sub>60</sub>	360	8	-61610.5	-61660.1

TABLE I. Convergence rate results for NLFEAST for various molecules (the molecular geometries are obtained from the experimental data<sup>23</sup>). The convergence criteria is satisfied for NLFEAST when the relative error on trace is below  $10^{-10}$ , which also provides very low non-linear relative residual  $\max_m (\|\mathbf{H}(\mathbf{x}_m)\mathbf{x}_m - E_m\mathbf{S}\mathbf{x}_m\|/\|\mathbf{H}(\mathbf{x}_m)\mathbf{x}_m\|)$ , typically here below  $10^{-8}$ . The total energy results using our cubic real-space FEM code are in good agreement with those obtained by NWChem<sup>22</sup> using the cc - pvqz basis from H<sub>2</sub> to C<sub>6</sub>H<sub>6</sub>, and the 6 - 31g\* basis for C<sub>8</sub>H<sub>10</sub>N<sub>4</sub>O<sub>2</sub> and C<sub>60</sub>.

iments presented in Figure 4 and Table I, the rate of convergence is expected to reach a maximum because all generated subspaces have been retained. Figure 5 shows the results of several numerical experiments that were performed for our selected three molecules using different sizes for the search subspace  $\hat{\mathbf{Q}}$ . From these results,




FIG. 5. Numerical experiments demonstrating the role of the number of most recent retained  $\mathbf{Q}^{(k)}$  subspaces for constructing the search subspace  $\hat{\mathbf{Q}} = \{\mathbf{Q}^{(k)}, \mathbf{Q}^{(k-1)}, \mathbf{Q}^{(k-2)}, \dots\}$  used by NLFEAST. The convergence rate depends on this number that here varies between one and “infinity” (i.e. where the search subspace keeps increasing until convergence is reached).

one observe that the rate of convergence of NLFEAST is directly related to the number of subspaces that are retained, but a high rate of convergence can still be obtained by retaining only three to four of the most recent

subspaces. By retaining a limited number of subspaces, it becomes also possible to increase the size of  $M_0$  to consider a larger number of electrons for a given search interval without resorting yet to explicit parallelism. In the case of Benzene and Caffeine, one also notes that the algorithm does not appear to converge at all when only a single  $\mathbf{Q}^{(k)}$  subspace is used, this subspace being the one that was most recently generated. Although this is not necessarily typical and the solution may eventually converge using a more efficient approach for solving the non-linear reduced problem (e.g. using a larger number of inner iterations for the FEAST-DIIS problem which has been fixed to three in our simulations), it highlights the importance of extending the search subspace size for the success of this algorithm.

### C. Convergence Rate and Contour Integration Accuracy

At each iteration of NLFEAST, the approximate subspace solution is improved through multiplication by the density matrix of the most current Hamiltonian (see equation 1). This step, Step 2 in Figure 3, is accomplished by performing a numerical contour integration of the Green’s function multiplied by the approximate solution. In practice, a  $n$ -point Gauss quadrature can efficiently be used here, which involves summing the solutions of  $n$  separate linear systems. For all the results




FIG. 6. Experiments demonstrating the effect on convergence of the accuracy of the contour integration in NLFEAST. Different curves represent the convergence for different numbers of Gauss points used in the numerical contour integration (Step 2 in Figure 3). For all these results, the search subspace  $\hat{\mathbf{Q}}$  keeps increasing indefinitely until convergence.

presented so far, 16 Gauss contour points were used to perform the quadrature. Figure 6 presents a new set of convergence results for our three selected molecules, while considering the variation of the number of Gauss points from as low as 4 to as high as 48. Although increas-



ing the number of Gauss points can help to improve the convergence rate, the effect here is not quite as dramatic as the effect of increasing the subspace size as discussed previously. We note here very similar convergence behaviors between NLFEAST and the FEAST algorithm for the linear problem. This latter admits a mathematical convergence proof that shows that the actual convergence rate depends on both the accuracy of the contour integration and the size of the search subspace<sup>15</sup>.

#### D. Robustness and initial guess

Like other schemes for solving nonlinear eigenvector equations, NLFEAST requires an initial guess for the density  $\mathbf{n}$ . Clearly, a good initial guess can provide faster convergence, but for NLFEAST the quality of the initial guess is not important as far as achieving convergence is concerned. Unlike other means of performing SCF iterations, our algorithm is capable of achieving convergence even when given an extremely poor initial guess including the extreme case of no initial guess at all (i.e.  $\mathbf{n} = 0$ ). Figure 7 shows the results of a number of numerical experiments wherein the initial guess for the electron density was set to zero. Each experiment eventually resulted in convergence. The algorithm’s performance when given




FIG. 7. Results of numerical experiments where the initial guess for the electron density was set to zero,  $\mathbf{n} = 0$ . The maximum number of subspace to form  $\hat{\mathbf{Q}}$  has been set equal to 10 for all molecules but  $\text{C}_8\text{H}_{10}\text{N}_4\text{O}_2$  and  $\text{C}_{60}$  where it was set equal to 4. The convergence is reached for an error on the trace smaller than  $10^{-9}$

such a poor initial guess is partly related to the size of the system; more electrons typically means that a larger number of iterations is required before convergence is reached. Performance also appears to depend on the particular molecule under consideration. In Figure 7,  $\text{CH}_4$  requires a larger number of iterations to reach a high level of convergence, despite being the second smallest system

(in terms of number of electrons). It is likely that, for some molecules, the algorithm converges towards a local energy minimum before ultimately finding its way to the global minimum, which results in its progress being more delayed than it would be with a molecule where such a detour does not occur. The purpose of Figure 7 is just to illustrate the robustness of the algorithm by considering an extreme (academic) case. As shown in Table I, a conventional initial guess has, so far, guaranteed low residual convergence in less than 10 FEAST iterations for all the molecules that we have experimented with.

#### V. CONCLUSION AND FUTURE WORK

A new eigensolver-based strategy, named NLFEAST, is derived from a generalization of the FEAST eigenvalue algorithm for solving the non-linear eigenvector problem i.e.  $H(\{\psi\})\psi = E\psi$  such as the one arising from the DFT/Kohn-Sham electronic structure model. By providing a fundamental and practical numerical solution for addressing the non-linearity with the occupied eigenvectors, NLFEAST offers a very efficient and robust alternative to the traditional self-consistent procedure using density-mixing schemes such as Pulay-DIIS. Several numerical experiments using the DFT/Kohn-Sham/LDA model, have demonstrated the significant potential of NLFEAST to outperform the traditional SCF mixing techniques in terms of convergence rate, robustness, and numerical operation counts. Strictly speaking, NLFEAST should not be considered as a direct competitor to current methods and it can rather be seen as an alternative formulation on how to handle the non-linearity. Indeed, the resulting non-linear reduced system can, in turn, be addressed using any SCF procedures. In our simulations, the reduced system did not need also to be solved accurately since the overall convergence of our scheme benefits from the robustness of the FEAST subspace iterations (which also limits the number of inner SCF iterations). This feature could potentially be further exploited while considering Hartree-Fock or DFT hybrid models where the construction of the Hamiltonian system is more involved than solving a single Poisson equation.

We note that a practical implementation of the technique can be achieved effectively using the FEAST solver package<sup>24,25</sup>. As such, the migration of electronic structure codes making use of SCF mixing schemes, could proceed within two steps: (i) integration of the FEAST package as the main linear eigenvalue solver, (ii) reverse-order the SCF-process as it was stated in Section III. Since the FEAST solver is reverse communication ready, it provides the flexibility to build a customized NLFEAST with specific computational modules for the problem at hand (independently of the physical and discretization models).

In this paper, the applicability of NLFEAST has also been demonstrated for a given search interval which can include several hundred states. In order to go beyond

the current performance capabilities of NLFEAST for addressing much larger molecular systems, including nanostructures of current technological interest, explicit parallelism will become necessary. The parallel treatment of NLFEAST should directly benefit from the intrinsic parallel capability of the FEAST algorithm (including the three level of parallelism mentioned in Section II). For the all-electron model, in particular, the algorithm can act on different energy ranges (with no overlap for the contour integrations), in order to capture core or valence electrons independently. For example, the results in Figure 8 clearly illustrate the various core and valence regions within the energy spectrum distribution for a given molecular configuration.




FIG. 8. Plots of the density of states for buckminsterfullerene ( $C_{60}$ ) and caffeine ( $C_8H_{10}N_4O_2$ ). Our all-electron code is able to capture both the valence states and the core states of each type of atom. From those results, we can clearly identify the low energy core regions of Oxygen, Nitrogen, and Carbon.

With many search intervals present in the simulation, however, one would also need to address the parallel scalability for solving the non-linear reduced system. Using the linear FEAST within a traditional SCF process, there is no issue, since the number of reduced systems is equal to the number of intervals. Using NLFEAST, however, the situation becomes less straightforward since the construction of the non-linear reduced system depends (in principle) on the solutions of the contour integrations for

all search intervals. The development of new direct or iterative schemes for solving the resulting reduced eigenvalue problem that scale linearly with the number of search intervals will need to be addressed in future research.

## ACKNOWLEDGMENTS

The authors wish to acknowledge helpful discussions with Prof. Joseph Jerome (Northwestern University, Evanston IL), Prof. Ahmed Sameh (Purdue University, West-Lafayette IN), and Dr. Ping Tak Peter Tang (Intel Corporation). This material is supported by NSF under Grant #ECCS-0846457 and Intel Corporation.

- <sup>1</sup>R. M. Martin, *Electronic Structure: Basic Theory and Practical Methods* (Cambridge University Press, 2004).
- <sup>2</sup>P. Pulay, *Chem. Phys. Let.* **2**, 393 (1980).
- <sup>3</sup>V. Eyert, *J. Comput. Phys.* **124**, 271 (1996).
- <sup>4</sup>D. Bowler and M. Gillan, *Chemical Physics Letters* **325**, 473 (2000).
- <sup>5</sup>T. Rohwedder and R. Schneider, *J. Math. Chem.* **49**, 1889 (2011).
- <sup>6</sup>A. Garza and G. Scuseria, *J. Chem. Phys.* **137**, 054110 (2012).
- <sup>7</sup>K. Kudin, G. Scuseria, and E. Cancès, *J. Chem. Phys.* **116**, 8255 (2002).
- <sup>8</sup>S. Høst, J. Olsen, B. Jansík, L. Thøgersen, P. Jørgensen, and T. Helgaker, *J. Chem. Phys.* **129**, 124106 (2008).
- <sup>9</sup>X. Hu and W. Yang, *J. Chem. Phys.* **132**, 054109 (2010).
- <sup>10</sup>Y. Wang, C. Yam, Y. Chen, and G. Chen, *J. Chem. Phys.* **134**, 241103 (2011).
- <sup>11</sup>K. Baarman and L. VandeVondele, *J. Chem. Phys.* **134**, 244104 (2011).
- <sup>12</sup>K. Baarman, T. Eirola, and V. Havu, *J. Chem. Phys.* **134**, 134109 (2011).
- <sup>13</sup>C. Yang, W. Gao, and J. C. Meza, *SIAM J. Matrix Anal. Appl.* **30**, 1773 (2009).
- <sup>14</sup>E. Polizzi, *Phys. Rev. B* **79**, 115112 (2009).
- <sup>15</sup>E. Polizzi and P. Tang, (2012), in preparation.
- <sup>16</sup>Z. Chen and E. Polizzi, *Phys. Rev. B* **82**, 205410 (2010).
- <sup>17</sup>S. Laux, *Phys. Rev. B* **86**, 075103 (2012).
- <sup>18</sup>C. Yang, J. C. Meza, and L.-W. Wang, *J. Comput. Phys.* **217**, 709 (2006).
- <sup>19</sup>J. Jerome, P. Sievert, L. Ye, I. Kim, and A. Freeman, *J. Comput. Elec.* **6**, 349 (2007).
- <sup>20</sup>A. Sameh, (2012), private discussion.
- <sup>21</sup>A. Levin, D. Zhang, and E. Polizzi, *Comput. Phys. Comm.* **183**, 2370 (2012).
- <sup>22</sup>M. Valieva et al., *Comput. Phys. Comm.* **181**, 1477 (2010).
- <sup>23</sup>“Computational chemistry comparison and benchmark database,” (2011), <http://cccbdb.nist.gov/>.
- <sup>24</sup>“Feast solver,” (2009), <http://www.ecs.umass.edu/~polizzi/feast/>.
- <sup>25</sup>E. Polizzi, “A high-performance numerical library for solving eigenvalue problems,” (2012), [arxiv.org/abs/1203.4031](http://arxiv.org/abs/1203.4031).